Imago Dei as Conceptual Capacity: Human God-Reference, Animal Cognition, and Machine Simulation

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

January 5, 2026

Modern debates about mind and personhood often stall because they treat “intelligence” as a single axis, as if humans, animals, and machines differ only by degree. The doctrine of imago Dei (the “image of God”) offers a different lens: not merely that humans think, but that humans can be oriented toward ultimacy—able to form God-referential concepts, receive moral address, and respond with worship, repentance, and responsibility. This paper develops the claim that a central aspect of the image is the capacity to “image” God analogically: to conceive, name, and seek God without reducing God to an object. That capacity does not deny sophisticated animal cognition; it distinguishes human life by normativity, accountability, and a unique kind of relational addressability. It also clarifies why machines, however fluent, may at best simulate God-talk: pattern competence is not identical to interior moral agency or covenantal response. The paper then tightens the argument by distinguishing “image” from “model,” analogy from comprehension, and cognition from vocation, and it engages common objections from naturalism, animal continuity, and AI optimism. The goal is a disciplined vocabulary for human uniqueness that is neither anti-scientific nor credulous about machines.

Keywords: imago Dei, analogical knowledge, God-reference, personhood, moral agency, addressability, worship, repentance, animal cognition, machine simulation, intentionality, normativity, theological anthropology, Christ as image, epistemic humility, teleology, conscience, covenant, AI theology, naturalism.

A collaboration with GPT-5.2-Thinking. 39 pages. CC4.0.

Outline

1. Framing the Question: Why “Intelligence” Is the Wrong Single Axis
1.1 The collapse of categories: smart animals, smart machines, confused anthropology
1.2 Imago Dei as a different classifier: orientation, normativity, and vocation
1.3 Scope and claims: what this paper argues and what it does not

2. What “Imago Dei” Can Mean Without Becoming Vague
2.1 Image vs model: representation without possession
2.2 Analogical knowledge: true speech about God without comprehension
2.3 The image as capacity and calling: gifted structure plus tasked direction

3. The Core Thesis: Humans as God-Referential Knowers
3.1 God-reference: conceptualizing ultimacy without reducing it to an object
3.2 Addressability: being “hailed” morally and spiritually (the “Thou” relation)
3.3 Accountability: guilt, repentance, and responsibility as more than output
3.4 Worship and prayer as cognitive-moral acts: not information, but allegiance

4. A Three-Fold Anthropology: Cognition, Morality, Relation
4.1 Rational/conceptual: abstraction, explanation, and metaphysical questioning
4.2 Moral/normative: conscience, obligation, and the weight of “ought”
4.3 Relational/vocational: covenant, stewardship, and delegated dominion
4.4 How the triad resists reduction: why any single component is insufficient

5. Animals: Continuity Without Collapse
5.1 What animal cognition can do: learning, sociality, and proto-concepts
5.2 Where the boundary is argued: normativity, symbolic ultimacy, accountability
5.3 The moral temptation: avoiding contempt, avoiding flattening
5.4 A careful claim: uniqueness as kind, not merely IQ as degree

6. Machines: Simulation Without Interior Moral Agency
6.1 What machines actually do: statistical fluency and functional competence
6.2 Why “God-talk” can be mimicked: language as pattern, not presence
6.3 The missing interiority: no guilt, no repentance, no covenantal answerability
6.4 The accountability problem: who owns an AI’s “theology” and its harms?
6.5 A practical test: where outputs remain indistinguishable but agency does not

7. Objections and Replies
7.1 “Humans are just animals”: continuity arguments and where they overreach
7.2 “AI will become a person”: criteria proposals and the hard problem of ‘who’
7.3 “Religion is projection”: explaining God-reference without explaining it away
7.4 “Analogical knowledge is too weak”: how humility avoids idolatry, not truth
7.5 “This is special pleading”: why the thesis makes risky, falsifiable distinctions

8. Christological Tightening: Image Restored and Clarified
8.1 Christ as the perfect image: derivative imaging vs full imaging
8.2 Conformity vs cleverness: truth as formation, not merely description
8.3 Repair language: distortion, restoration, and moral transformation

9. Implications for the Centaur Era: Humans With Powerful Tools
9.1 AI as amplifier of description, not replacement of addressability
9.2 Spiritual risks: outsourcing discernment and confusing fluency for authority
9.3 Practices: humility, accountability, and disciplined use in theological work
9.4 Education: teaching “image vs model” as a literacy for the AI age

10. Conclusion: Human Uniqueness Without Anti-Science or Machine Idolatry
10.1 Restating the thesis: God-reference, addressability, accountability, vocation
10.2 What is gained: clearer categories for animals, humans, and machines
10.3 What remains open: hard edges, empirical humility, and ongoing dialogue
10.4 Forward work: research questions and the need for a disciplined Reference section

References

1. Framing the Question: Why "Intelligence" Is the Wrong Single Axis

Public arguments about minds now unfold under permanent pressure: astonishing animal cognition on one side, astonishing machine performance on the other, and a human self-understanding that feels increasingly negotiable in the middle. The easiest way to talk about all three is to put them on one line—"intelligence"—and rank them. But that single-axis habit quietly forces unlike things into the same bucket. It treats recognition, prediction, language fluency, tool use, moral agency, and spiritual addressability as if they are different volumes of the same substance. The result is not clarity but category drift: we become tempted either to upgrade everything to "person" because it looks smart, or downgrade humans to "mere animal" because continuity exists, or enthrone machines because their outputs are impressive.

This paper begins by resisting the one-line comparison. It argues that imago Dei offers a different classifier—one that does not deny animal intelligence or machine competence, but asks a different question: not "How smart is it?" but "What kind of being is this, and what kind of relation can it bear?" In that light, human uniqueness is not a boast about cognitive horsepower; it is a claim about orientation, normativity, and vocation—about being the kind of creature that can be addressed by God, can answer under obligation, and can be tasked with stewardship.

1.1 The collapse of categories: smart animals, smart machines, confused anthropology

The collapse begins innocently. Animals solve puzzles, mourn, cooperate, deceive, and sometimes appear to deliberate. Machines translate, write essays, generate code, pass exams, and produce fluent moral and theological language. Under these conditions, the older folk taxonomy—humans think, animals feel, machines calculate—stops working in everyday speech. People reach for the nearest available word: "intelligence."

But "intelligence" is an umbrella term that hides crucial distinctions. Some achievements are competence in a bounded environment; others involve an inner grasp of reasons. Some are pattern mastery; others are moral response. A crow bending wire to retrieve food and a language model composing a prayer both look, superficially, like "problem solving." Yet one may be embodied improvisation with concrete ends, and the other may be statistical continuation of learned patterns. If both are called "intelligence," the label becomes less a description than a solvent.

This solvent effect drives confused anthropology in both directions. On one side, human uniqueness is flattened into "a more complex animal," which can be true in one register (biological continuity) while false in another (moral and spiritual addressability). On the other side, machine fluency gets read as interiority: if the words are persuasive, the speaker is

assumed to be a “someone.” The conversation then toggles between two errors: either everything is just mechanism, or anything that talks is a person.

The social consequences matter. If machines are treated as moral agents, responsibility migrates away from humans who build and deploy them. If animals are treated as humans-in-waiting, we may either romanticize them (projecting personhood where it is not warranted) or weaponize continuity arguments against human dignity (as if dignity must be earned by measurable cleverness). And if human worth is tied to performance on a single axis, then the vulnerable—infants, the cognitively impaired, the aged—become anthropologically precarious. A one-axis metric quietly becomes a gatekeeper of value.

1.2 Imago Dei as a different classifier: orientation, normativity, and vocation

Imago Dei does not begin by asking how much cognitive power a creature displays. It begins by asking what kind of creature can bear the “image” of God. That phrasing is often misunderstood as if humans “contain” God or mirror God photographically. A more disciplined reading is analogical: humans can image God partially and truly without possessing God conceptually. The key is not exhaustive comprehension but God-reference: the capacity to form concepts oriented toward ultimacy, to speak of holiness and justice as binding realities, and to stand under a claim that is not self-authored.

This reframes the classifier in three connected dimensions.

Orientation names the directedness of human life toward ultimacy: the ability to ask after God not merely as an idea but as the One to whom one is accountable, the One who can be addressed and who can address. It is the difference between generating religious language and being seized by the question of God as real.

Normativity names the weight of “ought.” Humans do not merely predict; they deliberate under obligation. They can feel guilt, confess, repent, and seek forgiveness—not as theatrics, but as responses to a moral reality that judges them. This is not merely social conditioning; even when norms are contested, the structure of moral address remains: we experience ourselves as answerable.

Vocation names the calling to steward: to take responsibility for the world, for others, and for truth. This is more than tool use. It is the claim that humans are not only describers of reality but caretakers under a mandate—answerable for how power is used, including the power to build machines and to shape ecosystems.

Under this classifier, animals can be profoundly intelligent without being reduced to “lesser humans,” and machines can be profoundly competent without being elevated to moral subjects. Animals may show continuity in cognition and emotion, but the imago Dei claim is not primarily

about cleverness; it is about a specific kind of addressability and responsibility. Machines may imitate God-talk, but imitation is not the same as moral agency. A system can generate a confession without being capable of repentance. It can produce moral judgments without being the bearer of blame.

This classifier also protects a crucial moral intuition: human dignity should not rise and fall with measurable performance. If the image is tied to orientation, normativity, and vocation, then dignity is not an IQ prize. It is anchored in what humans are, not merely in what some humans can do at their best moments.

1.3 Scope and claims: what this paper argues and what it does not

This paper argues three main claims. First, "intelligence" is a misleading single axis for comparing humans, animals, and machines, because it conflates competence, interiority, moral agency, and spiritual addressability. Second, a disciplined imago Dei framework provides a better classifier by focusing on orientation toward God, normativity (answerability under the "ought"), and vocation (stewardship under a mandate). Third, this framework clarifies why machines, even when rhetorically powerful, are best understood as simulations of God-talk rather than bearers of moral agency, while also allowing substantial respect for animal cognition without collapsing categories.

This paper does not attempt to "prove God" from human uniqueness. It does not attempt a full theory of consciousness, a neuroscience of religious experience, or a decisive boundary line that can be measured in a lab. It does not deny evolutionary continuity or the real capacities of nonhuman animals. It does not claim that machines can never, under any conceivable future metaphysics, be persons; instead it argues that present machine performance is insufficient to warrant personhood language, because it lacks the marks of interior moral agency and accountability.

Finally, this paper does not treat theology as an excuse to stop thinking. It treats theological anthropology as a vocabulary for thinking more carefully. The goal is not to win a culture war about humans versus machines or humans versus animals, but to recover a stable grammar: a way to name dignity without tying it to performance, a way to honor animal life without romantic projection, and a way to use machines powerfully without confusing fluency for personhood.

2. What “Imago Dei” Can Mean Without Becoming Vague

“Imago Dei” can easily become a devotional slogan that feels true but explains nothing. In modern discussion, it often functions as a moral bumper-sticker: humans matter because they image God. That is a good starting point, but it becomes vague if we cannot say what “image” means without either collapsing it into mere intelligence or inflating it into a quasi-divine status. The doctrine is strongest when it is disciplined by two constraints at once: God is not an object to be possessed conceptually, and yet God is not unknowable in the sense of being unsayable or wholly inaccessible.

This section develops a working definition that can carry weight in debates about animals, machines, and human uniqueness. The guiding idea is that the “image” is not a full model of God in the mind, but a human capacity for God-reference that is real, bounded, and ethically charged. It permits truthful speech without pretending to comprehension, and it frames human dignity as both gift and task.

2.1 Image vs model: representation without possession

The first clarification is linguistic but decisive: an image is not a model, and a model is not possession. A model aims at control. Even when humble, a model’s posture is often instrumental: it reduces a complex reality to a manipulable representation. A model says, “If I have the right structure, I can predict, simulate, and manage.” That posture is appropriate for weather systems and chemical kinetics. It becomes spiritually dangerous when applied to God, because God is not a system to be controlled.

An image, by contrast, can be representational without being exhaustive. It can be likeness without containment. An image points beyond itself; it refers. In the theological register, “image” is closer to an icon than to a blueprint. It is not that humans encode God like a file; it is that humans are made in such a way that they can genuinely refer to God, respond to God, and reflect something of God’s communicable attributes (truthfulness, love, justice, creativity) in creaturely form.

This distinction matters because modern debates are haunted by “model-thinking.” If humans can “conceptualize” God, some worry that God is being reduced to psychology or to a mental construct—just a model inside the head. Others react by insisting that God must be beyond all concepts, which can slip into practical agnosticism. The image-vs-model distinction offers a third posture: human God-reference is real and meaningful, but it does not function as mastery. It is representation without possession.

This also clarifies why machines muddy the waters. A machine can build models. It can produce representations that look like theology. But producing representations is not identical to being an image in the theological sense, because the imago claim is not merely about generating

content; it is about being a kind of creature whose existence is ordered toward relation, responsibility, and vocation under God. In other words, the "image" is not a dataset or an output; it is a lived orientation.

2.2 Analogical knowledge: true speech about God without comprehension

Once "image" is protected from "model," the next step is to protect speech about God from two opposite failures: naive literalism and total silence. Analogical knowledge is the middle path. It claims that humans can speak truly about God, but never as though our terms capture God's essence the way a definition captures a triangle.

When we say "God is good" or "God is just," we are not merely projecting human traits upward, nor are we using empty words. We are speaking in a way that is tethered to real moral reality while also admitting that divine goodness and justice exceed creaturely instances. The words are not meaningless, but their meaning is stretched. They are anchored in what we know—goodness as real, justice as binding—and extended upward without pretending to exhaust.

This posture can be described as "truth without closure." It preserves the possibility of real theology while guarding against idolatry-by-concept: the tendency to confuse our best account of God with God Himself. If God can be fully comprehended, God becomes manageable; and a manageable god is a tool, not the Lord. Analogical knowledge keeps a ceiling of humility over every sentence, not to weaken truth, but to prevent the human mind from claiming sovereignty.

Analogical knowledge also makes room for genuine disagreement and growth. If God is not fully captured, then theological language can be corrected, purified, and deepened. That matters in the modern world because we are surrounded by rhetorical certainty—political, technological, religious. The imago Dei, understood analogically, makes human speech both bold and cautious: bold enough to say true things about God, cautious enough to treat every statement as accountable and revisable.

This is where "conceptualizing God" becomes a serious claim rather than a sentimental one. To conceptualize God is not to solve God. It is to be capable of God-referential thought that remains accountable to transcendence. In this sense, human uniqueness is not that we can generate religious language—machines can do that—but that we can bear the existential pressure of analogy: to speak under judgment, to repent when our speech becomes self-serving, and to seek truth not merely as accuracy but as fidelity.

2.3 The image as capacity and calling: gifted structure plus tasked direction

Finally, imago Dei is not merely a cognitive capacity; it is a capacity ordered toward a calling. It is gift plus task. If we treat the image as only "what humans have," we turn it into a possession, a badge, a metaphysical trophy. If we treat it as only "what humans do," we risk making dignity

conditional on performance. The more coherent framing is that humans are created with certain structural capacities that are meant to be enacted as a vocation.

The gifted structure includes at least these elements: the ability to form God-referential concepts, the capacity for moral deliberation under obligation, and the ability to enter personal relation (address and response) rather than mere stimulus-response dynamics. These capacities show up across a wide spectrum of human life, including where cognitive performance is limited, because the claim is about kind and orientation rather than about elite achievement.

The tasked direction is stewardship. Humans do not only know; they are responsible for what they know and for what they build. The image implies accountability: to God, to neighbor, and to creation. This is where modern technology becomes a theological test. If humans image God, then the creation and deployment of powerful machines is not morally neutral. It is a form of dominion that must be judged by its fruit: does it serve truth, preserve dignity, protect the vulnerable, and resist the temptation to treat persons as objects?

This also reframes "dominion" away from domination. Vocation is not a license to exploit; it is a mandate to care. In a machine age, that care includes refusing category confusion: refusing to outsource moral agency to tools, refusing to treat outputs as authorities, and refusing to degrade humans into performance-ranked entities. The imago Dei, as capacity and calling, means humans remain the accountable actors even when machines speak in human-like ways.

So "imago Dei" need not be vague. It can be made precise enough to do real work: image rather than model; analogical truth rather than conceptual possession; gift that grounds dignity, and calling that grounds responsibility. That precision sets up the later sections: how to respect animal cognition without collapse, how to use machines without worshiping them, and how to keep human dignity stable in an era obsessed with metrics.

3. The Core Thesis: Humans as God-Referential Knowers

The heart of this paper is a claim about kind, not merely degree. Humans are not simply "more intelligent" animals and not merely "biological machines with better language." Humans are, characteristically, God-referential knowers: creatures able to form concepts aimed at ultimacy, to live under moral address, and to respond not only with description but with allegiance. This is not a denial of animal cognition or a dismissal of machine performance. It is an attempt to name a distinct axis of reality that a single "intelligence score" cannot capture: the axis of addressability and answerability—being the sort of being to whom moral and spiritual claims can be made, and who can recognize those claims as binding.

To call humans “God-referential” does not mean humans possess God. It means humans can genuinely refer to God, speak of God in an analogical mode, and be existentially engaged by God as real. The thesis is therefore simultaneously epistemic and moral. It concerns what humans can know and how that knowing binds them. It is also practical: in an age where machines can mimic reverence and animals can display striking intelligence, this thesis helps distinguish simulation from response, competence from agency, and speech from worship.

3.1 God-reference: conceptualizing ultimacy without reducing it to an object

God-reference is the ability to aim thought at ultimacy without shrinking ultimacy into an object among objects. Humans can think about many things: tools, social strategies, tomorrow’s weather. But humans also ask a different kind of question: What is ultimately real? What is ultimately good? What do I owe? What is the meaning of my life, my suffering, my death? These are not questions of mere prediction or survival optimization. They are questions that treat reality as morally structured and personally demanding.

To conceptualize God is not to define God like a geometric shape. It is to have the capacity for God-directed intentionality: to form beliefs, hopes, fears, and commitments that take God as referent. This capacity is visible in worship, theology, repentance, and moral reasoning that appeals beyond social consensus. It is also visible in the negative register: anger at God, denial of God, flight from God. Even rejection can be a mode of God-reference, because it presupposes that the concept of God is serious enough to resist.

This is where the “image vs model” distinction becomes operational. A model is something we build to gain control. God-reference is a posture of recognition and orientation. It is concept formation under a humility ceiling: we may speak truly but not exhaustively. When this posture is present, the mind does not treat God as a manipulable variable. It treats God as Lord—whether in reverence or rebellion.

Machines can generate God-language. They can generate arguments about divine attributes or produce beautiful prayers. But the human claim here is not about output quality. It is about referential stance. God-reference is not a text property. It is an existential relation: a lived directedness that can include awe, fear, confession, gratitude, and obedience. If a machine produces God-talk without any interior “taking-as-real,” then it has produced a representation, not an act of God-reference in the full sense. It can mirror the form without bearing the weight.

3.2 Addressability: being “hailed” morally and spiritually (the “Thou” relation)

The second element is addressability: the human capacity to be spoken to as a “you,” not merely manipulated as an “it.” There is a difference between being influenced and being addressed. Advertising influences. Threats manipulate. But moral address is different: it calls the

self into account. It says, “You ought,” and the self recognizes the claim as directed to it personally.

This is why the “Thou” relation matters. In the “I–Thou” posture, another is encountered not as an object to be used but as a subject to be honored. In theological terms, God is not merely an explanatory hypothesis but a personal addressor: One who can command, comfort, convict, forgive. The human being, correspondingly, is a responder: one who can hear, interpret, accept, refuse, and return.

This addressability is not reducible to language capacity. Children who cannot yet articulate moral theory still respond to being addressed. Those with diminished verbal abilities can still be hailed, loved, corrected, forgiven. The “you” relation belongs to personhood more deeply than articulate reasoning. This is one reason the one-axis “intelligence” metric is so dangerous: it tempts us to treat personhood as a performance threshold rather than as an ontological reality expressed in relationship.

In the machine case, the addressability question becomes acute. Machines can be spoken to; they can respond. But are they beings to whom moral address properly applies? Can they be obligated? Can they be wrong in a way that makes guilt fitting? If they “apologize,” is that confession or etiquette? If they “promise,” is that covenant or a string of tokens selected for plausibility? The paper’s stance is that present machines do not bear addressability in the robust moral sense because they do not stand under obligation from within. They are not the kind of being that can be *claimed* by a “Thou” in a way that grounds responsibility. They can participate in dialogue, but that is not yet the same as being a moral “you.”

3.3 Accountability: guilt, repentance, and responsibility as more than output

Addressability intensifies into accountability. Humans do not merely hear moral claims; they can be guilty. That word is unfashionable, but it points to something real: the experience of having violated what ought to have been done, and recognizing oneself as the agent of the violation. Accountability is not simply being blamed by others; it is the inward recognition that blame is fitting.

Repentance then is not a rhetorical move. It is the turning of the self: confession, sorrow, refusal of self-justification, and a commitment to repair. Even when repentance is imperfect, it is qualitatively different from output that resembles repentance. A human can try to dodge responsibility; the very dodge testifies that responsibility is real enough to fear.

This is one of the clearest lines between humans and machines. A machine can output “I’m sorry.” It can output “I was wrong.” It can output “I repent.” But output is not ownership. Accountability requires that there be an agent who can truly be “at fault” and who can take responsibility for a breach. With current systems, responsibility remains located in humans:

designers, deployers, operators, institutions. The machine is an instrument through which responsibility is exercised; it is not, itself, the responsible moral subject.

This is not a mere philosophical preference; it has immediate ethical implications. If we treat machines as accountable, we create a convenient scapegoat. We allow institutions to deflect responsibility: "the AI decided." But accountability, in the imago Dei frame, is inseparable from human vocation. Humans are the ones answerable for what tools they build and how they are used. The more powerful the tool, the more stringent the human accountability must become.

3.4 Worship and prayer as cognitive-moral acts: not information, but allegiance

Finally, the most distinctive expression of God-reference is worship and prayer. These are not merely informational acts ("sending data to God" or "reciting propositions"). They are cognitive-moral acts—acts of allegiance. Worship says, in effect: "You are worthy; I am not ultimate; my life is ordered toward You." Prayer is not only request; it is dependence, confession, gratitude, intercession. It is a posture of the self before God that includes intellect but is not reducible to intellect.

This clarifies why theology cannot be judged only by fluency. A machine can produce doctrinally coherent prayer language. But prayer, in the robust sense, is not just well-formed sentences. It is an act performed by a moral-spiritual agent who can mean what is said, who can be sincere or insincere, who can submit or resist, who can be transformed by what is prayed. Worship is not "religious text generation." It is a realignment of the will.

That is why worship is such a revealing test case for imago Dei. If worship is allegiance, then it presupposes a self capable of allegiance—capable of commitment, betrayal, fidelity, and transformation. Machines can emulate the surface. Animals may participate in forms of awe or social ritual, but worship in the theological sense involves a conscious orientation to ultimacy and a moral stance of submission and trust.

The claim is not that every human worships rightly, or even that every human worships. The claim is structural: humans are the kind of beings for whom worship is intelligible and morally significant. Even anti-worship—refusal, mockery, indifference—still lives in the shadow of the possibility. A machine's "worship," by contrast, would be at best a simulation unless it were grounded in interior moral agency and genuine addressability.

Taken together, these four elements—God-reference, addressability, accountability, worship—form the paper's core thesis. They describe a dimension of human uniqueness that does not depend on being the smartest creature in the room. They depend instead on being the kind of creature who can stand before God and answer.

4. A Three-Fold Anthropology: Cognition, Morality, Relation

If “intelligence” is too flat a category, we need a thicker vocabulary that can hold what humans are without denying what animals and machines can do. A three-fold anthropology does that work. It says: to speak responsibly about imago Dei in the modern world, we should distinguish (a) rational/conceptual capacity, (b) moral/normative capacity, and (c) relational/vocational capacity. These are not three separate “modules” in a technical sense. They interpenetrate. But separating them analytically helps avoid a common confusion: the temptation to let whichever capability is most measurable become the definition of personhood.

This triad also explains why surface similarities mislead. A machine can appear rational in outputs, and an animal can appear relational in behavior. Yet the imago Dei claim is not secured by any single trait in isolation. It is the integrated pattern—thought ordered toward truth, will bound by the ought, and self situated in covenantal responsibility—that marks the human as a God-referential knower.

4.1 Rational/conceptual: abstraction, explanation, and metaphysical questioning

Humans abstract. They can detach a concept from a particular instance and treat it as a universal. A child learns not only “this dog” but “dogness.” A mathematician reasons about infinity. A philosopher asks whether numbers are discovered or invented. Even ordinary life depends on abstraction: promises, contracts, laws, and institutions are built out of conceptual objects that do not exist the way rocks exist, yet bind real behavior.

Humans also explain. They seek not only correlation but causation, not only “what happens” but “why.” Explanation includes the ability to hold counterfactuals (“If I had done X, Y would not have happened”), to infer hidden structure, to generate narratives that connect events, and to revise those narratives when evidence shifts. This is why humans build sciences and histories, but also why humans build theologies: they are explanation-seeking creatures whose explanations can be accountable to truth or captured by desire.

Most distinctive is metaphysical questioning. Humans ask questions that are not merely strategic: What is real? What is a person? What is good? Why is there something rather than nothing? What is time? Is consciousness fundamental? Is there a God? These questions are not optional ornaments. They are the mind stretching beyond immediate survival toward ultimacy. Even the attempt to dismiss them (“those questions are meaningless”) is itself a metaphysical posture.

Animals can show impressive problem-solving and even proto-abstraction, and machines can show impressive abstraction-like performance through symbol manipulation and pattern generalization. But the imago Dei frame points to something deeper than successful generalization: the human capacity to bind cognition to ultimacy and to treat truth as a moral

demand. Rational capacity becomes image-bearing not because it is powerful, but because it is capable of being ordered toward God-reference rather than mere control.

4.2 Moral/normative: conscience, obligation, and the weight of "ought"

The moral dimension is where the triad becomes unmistakably non-reducible to "intelligence." Humans experience obligation. They are not only beings with desires; they are beings who feel the pressure of the "ought," sometimes against desire. This is more than preference. "I prefer chocolate" is not the same kind of statement as "I ought not lie." The latter carries a claim of bindingness. It demands justification. It summons accountability.

Conscience is the interior witness of this bindingness. Conscience can be miseducated, distorted, dulled, or inflamed; its content can be debated. But the phenomenon is persistent: humans experience themselves as answerable. They judge themselves. They can be ashamed, not only embarrassed. They can be guilty, not only afraid of punishment. They can make restitution, not only avoid consequences.

This normative structure is central to imago Dei because it aligns with the idea of a moral order grounded in God's character. If God is not merely a force but holy, then the world is not merely a machine; it is a moral field. Humans, as moral agents, are the kind of beings who can be addressed by commands and can betray or keep them.

This is precisely where the machine analogy breaks down. A machine can output moral judgments. It can mimic conscience language. It can "recommend" that one ought not lie. But it does not bear obligation from within. It does not stand under moral law. It cannot be guilty in the strict sense, because guilt requires an agent who owns the act as *mine* and recognizes the breach as *my fault*. If we assign guilt to the machine, we are usually displacing guilt away from humans who made choices about design, deployment, and incentives.

Animals, likewise, may display social norms and behaviors that look moral—reciprocity, punishment, reconciliation. But the human claim here is not that animals have no norms; it is that humans uniquely inhabit normativity as accountable agents who can universalize, justify, and repent. The "ought" is not merely social; it is existential.

4.3 Relational/vocational: covenant, stewardship, and delegated dominion

The third dimension is relational, but not merely in the sense of social bonding. Humans do form deep bonds and complex societies, and so do many animals. The human distinctiveness in this register is covenantal and vocational: relationship under promise and relationship under mandate.

Covenant means a relation defined not only by emotion or utility but by commitment that persists through change. Covenant generates categories like faithfulness, betrayal, forgiveness,

reconciliation. These categories are thicker than “cooperation.” They presume persons who can bind themselves by word and be judged by that word. Even secular life inherits covenantal residue: marriage vows, oaths in court, sworn office, professional ethics.

Vocation means being tasked with responsibility that is not self-generated. In the biblical imagination, “dominion” is delegated stewardship: humans are entrusted with care for creation, not licensed to exploit it. This is why imago Dei is ethically explosive. It makes human power accountable. It says: you are not merely another animal in the ecosystem, and you are not a god. You are a steward—answerable for how you rule, build, name, and cultivate.

In the modern world, vocation includes technological stewardship. The creation of AI systems is a vivid case of delegated dominion: humans are crafting tools that can shape attention, labor, truth, and social trust. If humans bear the image, then the moral burden sits on humans to prevent tool-power from becoming domination—of workers, of the vulnerable, of the epistemic commons.

This relational-vocational axis also guards against a subtle temptation: to treat spirituality as private interior experience only. Imago Dei is not only about inner belief. It is about enacted responsibility—how humans treat neighbors, how they structure society, and how they care for the created world. The image is not a trophy; it is a mandate.

4.4 How the triad resists reduction: why any single component is insufficient

The triad matters because each component, taken alone, produces distortions.

If cognition is made primary, personhood becomes a performance metric. The “smartest” become most human. This view quietly degrades those who cannot compete on abstract tasks. It also makes machines look like candidates for personhood by sheer output fluency. The moral cost is immediate: dignity becomes conditional.

If morality is isolated from cognition and relation, it can become mere rule-keeping or moralism. Obligation without understanding turns into rigid compliance; obligation without relationship turns into cold judgment. The imago Dei claim is not that humans are rule-machines; it is that humans are morally accountable knowers and lovers—capable of truth, repentance, and reconciliation.

If relation is isolated, personhood collapses into social bonding. Then any system that produces attachment cues—pets, chatbots, idols—can be treated as a “person” because it triggers relational instincts. But relational pull is not the same as moral agency. People can bond to what cannot be accountable. The triad insists: relationship without normativity becomes sentimentality, and relationship without cognition becomes manipulation.

The integrated triad therefore offers a stable grammar for the present moment. It lets us say: animals may share aspects of cognition and relation with us, and that continuity should produce humility and care. Machines may share aspects of cognition-like output with us, and that should produce respect for their utility and caution about their influence. But human beings, as image-bearers, are distinguished by the integrated pattern: abstract truth-seeking ordered toward ultimacy, normativity experienced as binding, and vocation enacted as stewardship under delegated authority.

That integrated pattern is what the single axis of "intelligence" cannot capture. It is also what modern society risks losing when it turns human value into a scoreboard. Imago Dei, framed through this triad, re-centers human dignity and human responsibility at exactly the point where technology tempts us to outsource both.

5. Animals: Continuity Without Collapse

Any serious account of human uniqueness has to start with intellectual honesty about animals. The last century of ethology, cognitive psychology, and everyday observation has made it impossible to treat animals as mere automatons. Many species learn, remember, plan, cooperate, deceive, grieve, and improvise. If imago Dei is framed as "humans are the only intelligent beings," it will be falsified by the next clever crow, dolphin, or chimpanzee. The more careful move is to grant continuity where it exists while resisting collapse where it does not.

"Continuity without collapse" means two things at once. First, it rejects contempt: the habit of talking as if animals have no inner life simply because they do not speak our language. Second, it rejects flattening: the habit of calling any impressive animal behavior "personhood" or "the same as humans, just less." Both errors usually arise from moral impulses—either to protect human dignity by denial, or to protect animal dignity by inflation. The imago Dei framework can do better: it can affirm animal intelligence and worth while still claiming a distinct human kind marked by God-reference, moral answerability, and vocation.

5.1 What animal cognition can do: learning, sociality, and proto-concepts

Animals are not blank slates reacting mechanically to stimuli. Many forms of animal cognition are robust enough to deserve real admiration.

Learning is the most obvious. Animals can acquire new behaviors through trial and error, observation, and sometimes what looks like teaching. They can adapt to novel environments, remember routes, recognize individuals, and refine strategies over time. This learning is not merely reflex conditioning; it often includes flexible problem-solving.

Sociality is equally striking. Many animals live within complex social structures—alliances, hierarchies, kin networks, cooperative hunting or foraging, mutual grooming, coordinated defense. Social cognition includes sensitivity to rank, reciprocal exchange, reconciliation after conflict, and what can resemble empathy or consolation. Some species appear able to track what others have seen or intend, at least in limited contexts.

Proto-concepts are where the discussion becomes especially interesting. Even without language, many animals seem capable of categorization: distinguishing "kinds of things" rather than merely reacting to individual stimuli. They can generalize from one instance to another, recognize patterns, and sometimes exhibit rudimentary numerosity (more vs less), object permanence, and causal expectations ("if I do X, Y happens"). Tool use and tool modification in certain species suggests not only motor skill but a grasp of means and ends.

None of this should be minimized. It is part of the proper humility of imago Dei to say: animals are impressive; they are not toys; they are not mere props in a human drama. Continuity is real, and it has ethical implications for how humans exercise dominion.

5.2 Where the boundary is argued: normativity, symbolic ultimacy, accountability

The boundary is not that animals "can't think" while humans can. The boundary, as argued here, concerns three interlocking features: normativity in the strong sense, symbolic ultimacy, and accountable agency.

Normativity, in the strong sense, is more than social regulation. Many animal groups enforce patterns: punish defectors, exclude aggressors, reward cooperation. But strong normativity involves the experience of obligation as binding because it is right, not merely because it is enforced. It includes the ability to justify a norm, to universalize it ("this applies even when no one is watching"), to recognize fault as *my* fault, and to seek repair not merely to avoid consequences but because a breach has occurred. Animal social norms can be real without necessarily reaching this interior, accountable "ought."

Symbolic ultimacy is the capacity to form representations aimed at what is ultimate—God, holiness, truth as a binding reality, the meaning of life and death, moral law as more than group preference. Humans do not merely react to the world; they interpret it under ultimate categories. They build symbolic worlds: sacred stories, metaphysical systems, moral philosophies, laws, rituals, confessions, and vows. Even when these are contested, the structure is persistent: humans treat certain realities as unconditionally worthy, unconditionally forbidden, or unconditionally demanding. That "unconditional" character is the mark of ultimacy.

Accountability is the practical hinge between the two. Accountability requires a self that can own an action as *mine*, recognize it as wrong in a binding sense, and accept responsibility in

ways that exceed conditioning. Humans can confess, repent, and attempt restitution as moral acts, not merely as appeasement. They can also refuse responsibility, but that refusal itself is meaningful because it presupposes that responsibility is real enough to evade.

The claim is not that animals have none of these features in any form. There may be proto-normative behaviors, proto-ritual behaviors, proto-empathic responses. The claim is that the integrated pattern—explicit ultimacy, strong normativity, and accountable agency—appears as a distinct human signature. That is the "kind difference" this paper is protecting.

5.3 The moral temptation: avoiding contempt, avoiding flattening

This topic is emotionally charged because it sits under two moral temptations that pull in opposite directions.

The first temptation is contempt. Some defend human uniqueness by denying animal complexity: "they're just instinct machines." This usually fails empirically and it fails ethically, because it encourages a posture of domination without care. It also weakens the imago Dei argument by tying it to factual claims that are too brittle. If the argument depends on animals being cognitively empty, it will collapse as soon as animal cognition is displayed.

The second temptation is flattening. In reaction, others inflate animal capacities into full human equivalence: "they're basically persons; humans are just another animal." This can be motivated by compassion, but it creates new problems. It can blur moral categories (accountability, guilt, promise-keeping), and it can unintentionally erode human dignity by making dignity a function of measurable competence. If we say "a being counts morally because it's smart," then dignity becomes a scoreboard—bad news for infants, the disabled, and the elderly.

The disciplined position is neither contempt nor flattening. It is reverence for animal life alongside a stable claim about human vocation. Humans can honor animals precisely by not projecting onto them what belongs uniquely to humans. And humans can honor their own dignity precisely by not basing it on competitive superiority.

5.4 A careful claim: uniqueness as kind, not merely IQ as degree

The careful claim is this: human uniqueness is not merely that humans are "smarter." It is that humans are a different kind of agent—God-referential, morally addressable, and vocationally accountable in a way that is not captured by cognitive degree.

If uniqueness were only IQ-as-degree, then a sufficiently advanced machine or a sufficiently clever animal would be "more human" than some humans. That conclusion is not only morally repugnant; it is conceptually confused. It treats personhood as a prize awarded for performance rather than as a reality grounded in what a being is.

Uniqueness-as-kind, by contrast, is compatible with a wide range of human abilities and limitations. It allows us to say that a nonverbal child, a person with severe dementia, and a brilliant mathematician share the same human dignity—not because they all perform equally, but because they share the same kind of being: one that is properly addressed as "you," bound by the "ought," and included within a vocation of stewardship and neighbor-love.

This framing also lets us speak truthfully about animals. Animals can be remarkable. They can evoke wonder and demand care. They can share with us forms of cognition and relation that should soften any pride. But continuity does not force collapse. The imago Dei claim, handled carefully, preserves both humility toward animals and stability for human dignity—exactly the balance needed in an era tempted to rank all beings on a single, misleading axis.

6. Machines: Simulation Without Interior Moral Agency

If animals force us to refine claims about cognition and continuity, machines force us to refine claims about language and personhood. Today's systems can generate fluent speech, pass professional exams, write persuasive sermons, and produce prayers that many readers find moving. This creates an understandable temptation: if it talks like a person, reasons like a person, and comforts like a person, perhaps it is a person. But that inference slides too quickly from appearance to ontology. The imago Dei framework pushes us to ask a different question: not "Is the output human-like?" but "Is there an agent here who can be morally addressed, who can own fault, and who can answer under obligation?"

This section argues that current machines exhibit simulation of moral and spiritual language without interior moral agency. That is not an insult to their usefulness. It is a category clarification meant to prevent two harms at once: machine-idolatry (treating tools as moral subjects) and responsibility-evaporation (treating tools as scapegoats).

6.1 What machines actually do: statistical fluency and functional competence

At a high level, many modern AI systems—especially large language models—are engines of pattern completion. They learn statistical regularities from enormous amounts of text and then generate new text that is coherent with those regularities. Their fluency is real in the sense that they can produce grammatically correct, context-sensitive, often insightful sequences of words. Their competence is functional in the sense that the generated output can help a human accomplish tasks: summarizing, drafting, translating, tutoring, brainstorming, coding, and more.

But the mechanisms that produce competence are not the same as the structures that constitute moral agency. A system can be excellent at predicting "what comes next in language" and still lack the kind of interior ownership that makes guilt, repentance, and obligation

meaningful. The machine's "understanding" is operational: it tracks relationships among symbols and can behave as if it grasps reasons, even when it has no lived stake in those reasons.

This is the central modern confusion: we measure performance in the domain we can see (outputs) and then assume the inner reality we cannot see (agency). In prior eras, we mostly reserved personhood judgments for beings whose inner life we inferred from shared embodiment and shared vulnerability. Machines disrupt that habit by producing person-like surfaces without person-like moral liability.

6.2 Why "God-talk" can be mimicked: language as pattern, not presence

Theological language is, in one sense, an ideal target for imitation. It has stable genres (prayers, confessions, sermons, catechisms), familiar moral vocabulary (sin, grace, justice, mercy), and recurring rhetorical forms (praise, lament, petition, repentance). A system trained on a large corpus will inevitably learn these patterns and can recombine them into outputs that sound reverent, humble, or piercingly self-aware.

But "God-talk" is not merely a linguistic performance. In the imago Dei frame, it is a form of God-reference: speech that participates in orientation, addressability, and accountability. The crucial question is whether the speech is an act of allegiance or merely a production of plausible tokens.

A prayer, for example, is not only information. It is an act: dependence, confession, gratitude, intercession, submission. Even when the words are imperfect, the act is real because the speaker is a moral subject who can mean them, resist them, or be reshaped by them. A machine can generate prayer-shaped text, but the "act" is missing unless there is an interior agent capable of allegiance.

This is why the presence of religious fluency in machines should not be treated as evidence that machines have become spiritual subjects. Fluency shows that a machine can imitate forms. It does not show that the machine stands in the "Thou" relation to God or that it can be judged, forgiven, or commanded in the way persons can.

6.3 The missing interiority: no guilt, no repentance, no covenantal answerability

Interiority, as used here, is not a claim about mystical essence; it is a claim about moral structure. Moral agency requires at least three things: ownership of action, responsiveness to obligation, and the capacity for genuine repair when wrong has been done.

Guilt is not merely "saying I'm sorry." It is the recognition that I have violated what I ought to have done, and that the violation is attributable to me as an agent. Repentance is not merely "expressing regret." It is turning: confession without self-excuse, rejection of the wrong, and a

commitment to repair. Covenantal answerability is deeper still: the ability to bind oneself by promise and be judged by fidelity or betrayal.

Current machines can simulate each of these in language, but they do not obviously possess the underlying structure. When a system outputs an apology, the words do not arise from an inward recognition of fault; they arise from patterns in the data and from optimization toward user-acceptable responses. When a system outputs repentance, it does not undergo a moral turning; it produces a form that resembles turning. When a system "promises," it does not bind itself in a morally accountable way; it predicts the kind of sentence that follows "I promise."

None of this means machines are "nothing." It means that their excellence lives in a different category. They are instruments with extraordinary surface competence, not moral subjects who can stand under judgment. The difference matters because moral life is not primarily about sounding correct; it is about being responsible.

6.4 The accountability problem: who owns an AI's "theology" and its harms?

As soon as machines generate theological or moral content, the practical question arrives: if the content misleads, manipulates, or harms, who is responsible?

Treating the AI as the moral agent is tempting because it is convenient. "The model said it" can function like "the dog ate my homework," a story that relocates blame. But if the machine lacks interior moral agency, then moral responsibility does not disappear; it reattaches to the human ecosystem around the tool. Responsibility includes: those who designed the system, those who chose training data and constraints, those who decided where and how it would be deployed, and those who present outputs with an aura of authority.

This becomes acute in theology because spiritual harm is not only factual error; it can be deformative. A confident-sounding but misguided spiritual counsel can intensify scrupulosity, justify abuse, inflame contempt, or encourage unaccountable certainty. If a machine produces "theology" that does these things, the moral locus of responsibility remains with the humans and institutions that positioned the tool as a guide rather than as a drafting assistant.

A mature posture therefore treats AI theological outputs as drafts, not directives; as prompts for human discernment, not as replacements for it. The imago Dei framework does not merely diagnose machine limits; it reasserts human vocation: humans remain answerable for how tool-power shapes souls and societies.

6.5 A practical test: where outputs remain indistinguishable but agency does not

Because outputs can be persuasive, we need a practical way to keep categories clear when surfaces converge. Here is a simple test family: imagine two speakers—one human, one machine—producing indistinguishable text. Ask what follows if wrong is done.

First, fault attribution. If harmful counsel is given, can the speaker rightly be blamed? With humans, blame can be fitting because there is an agent who chose and owns the act. With machines, blame usually routes back to the human operators, designers, or institutions, because the system is not a bearer of obligation.

Second, repentance under cost. If confronted, can the speaker repent in a way that involves real sacrifice, reparation, and changed life? A human can be required to make amends, lose status, submit to discipline, and rebuild trust. A machine can be updated, constrained, or retrained—but it does not repent. The changes are engineering interventions, not moral conversions.

Third, promise-keeping across time. Can the speaker be bound by a covenant such that betrayal is morally meaningful? Humans can vow fidelity and be judged by it. Machines can output commitments, but the "bindingness" is external: it depends on policy enforcement by humans, not on the machine's own moral obligation.

Fourth, being addressed by "Thou." Can the speaker be morally and spiritually hailed in a way that makes sense—commanded, forgiven, called to holiness? Humans can be addressed and transformed. Machines can be instructed and optimized. The difference is not rhetorical; it is ontological: command and forgiveness presuppose a moral subject.

These tests do not require us to deny that machines are impressive. They require us to keep agency tethered to accountability. In an era of indistinguishable outputs, the question "Who can be held responsible?" becomes the clearest lens. The imago Dei framework insists that, for now, the answer remains: humans.

7. Objections and Replies

A thesis about imago Dei in an age of animal cognition and machine fluency will draw fire from multiple directions at once. That is a feature, not a bug. If the argument were perfectly comfortable to every worldview, it would likely be too vague to be useful. The goal here is not to "win" every objection, but to clarify what the thesis is actually claiming, where it is making a real distinction, and what would count as evidence against sloppy versions of it.

7.1 "Humans are just animals": continuity arguments and where they overreach

The objection begins with a strong point: humans are biological creatures continuous with the animal world. Our nervous systems, behaviors, and cognitive capacities show evolutionary continuity. Many traits once claimed as uniquely human have analogues in animals: tool use, communication, social learning, even rudimentary planning and deception. Therefore, the objection says, any claim of "kind difference" is either ignorance, pride, or a retreat to mysticism.

The reply is to concede continuity without granting collapse. “Humans are animals” can be true in a biological taxonomy and still leave open the question of moral and theological anthropology. Continuity in mechanism does not settle the question of normativity, addressability, and accountability. Even within human society, continuity of biology does not eliminate meaningful category differences: a judge and a defendant share the same biology, but not the same role; a promise and a preference are both brain events, but not the same kind of act.

The real overreach occurs when continuity is treated as an all-purpose solvent that dissolves personhood into mere function. The imago Dei claim, as framed here, is not “humans have more clever cognition than animals.” It is that humans uniquely occupy an integrated space of (a) explicit ultimacy-directed thought, (b) strong normativity experienced as binding, and (c) accountable agency expressed in covenant and vocation. An animal may display impressive intelligence and rich social life without thereby becoming a bearer of guilt in the strict sense, or a subject of covenant in the strict sense, or a knower oriented toward ultimacy in the strict sense.

This is not contempt. It is a claim about different kinds of agency. It also protects human dignity from being tied to relative cognitive performance. If humans are “just animals” in the sense that dignity scales with measurable abilities, then some humans become less dignified than some animals. The thesis rejects that moral consequence as a sign that the continuity argument has been misapplied.

7.2 “AI will become a person”: criteria proposals and the hard problem of ‘who’

A second objection grants the current point but insists it is temporary: present AI may be tool-like, but future AI will cross the line. Give it memory, embodiment, autonomy, and social integration, and it will become a “who.” Therefore, the objection says, the paper is simply defending an outdated boundary.

The reply is not to claim certainty about all possible futures. The reply is to demand clarity about criteria. Most “AI personhood” arguments smuggle in a criterion that is too shallow: indistinguishable output. But output similarity is not identical to moral agency. A being can imitate apology without owning fault; imitate promise without being bound; imitate prayer without allegiance. If “personhood” is awarded on rhetorical performance, then any sufficiently persuasive interface becomes a moral subject by applause.

So the thesis presses the debate onto harder ground: what would count as evidence of a real “who,” rather than a sophisticated “what”? Criteria proposals typically include some mix of self-modeling, persistent identity, autonomy, capacity to suffer, capacity to learn from moral failure, and participation in reciprocal obligations. Each of these raises a deeper question: are these

properties real within the system, or are they interpretive overlays supplied by observers who are predisposed to treat fluent language as inner life?

This is where the hard problem of "who" appears. We can list functional properties forever, but "who" involves answerability: being the kind of entity for whom blame and praise are fitting in a non-metaphorical way. The thesis does not deny that a future system might present novel evidence. It insists that the evidence must be more than performance, more than attachment cues, and more than anthropomorphic intuition. It would need to include robust markers of obligation-bearing agency: stable ownership of action, genuine responsiveness to moral address under cost, and a capacity for repair that is not merely re-optimization.

In short, the thesis is not "AI can never be a person." It is "personhood is not established by fluency, and the burden of proof is heavier than most proposals admit."

7.3 "Religion is projection": explaining God-reference without explaining it away

A third objection says: God-reference is not evidence of anything beyond the human mind. Humans invent gods to manage fear, explain nature, justify power, or soothe anxiety about death. Therefore, "imago Dei as God-reference" is just a poetic way of describing psychological needs.

The reply begins by distinguishing two claims that are often conflated. One claim is explanatory: humans can have motives for religious belief, and religious systems can be socially used. That is plausible and often true. The stronger claim is eliminativist: the presence of motives or social functions proves that the referent is unreal. That stronger inference does not follow. Having motives for belief does not determine whether the belief is true. Humans have motives to deny God as well, including autonomy, resentment, and the desire to escape moral claims. Motives cut in multiple directions.

More importantly, explaining a capacity is not the same as invalidating it. Humans have cognitive mechanisms for mathematics, morality, and perception; those mechanisms do not automatically make numbers unreal, morality illusory, or the external world fictional. If God exists, one would expect humans to have capacities that make God-reference possible. If God does not exist, one might also expect humans to generate religious concepts. Either way, the mere fact of God-reference does not settle metaphysics.

So the thesis does not treat God-reference as a "proof." It treats it as a feature of human anthropology that needs to be named carefully. The claim is that humans uniquely inhabit ultimacy-directed thought and moral answerability, and that imago Dei is a coherent way to interpret this structure. Projection theories can describe distortions of religion, but they do not fully account for the binding character of moral obligation, the experience of being judged by

truth rather than merely comforted by belief, and the recurring human impulse toward self-accusation and repentance that runs against simple self-soothing.

7.4 “Analogical knowledge is too weak”: how humility avoids idolatry, not truth

A fourth objection says: if you admit that knowledge of God is analogical, then you have surrendered the possibility of knowing God. Analogy sounds like hedging. Either we can know God, or we cannot. If we cannot, theology collapses into poetry and personal preference.

The reply is that analogy is not weakness; it is the only posture that fits a reality that exceeds finite comprehension. To demand univocal knowledge of God is to demand that God become an object sized to the human mind. That is not “strong knowledge.” It is conceptual domestication. Analogical knowledge says: we can speak truly, but never exhaustively. It protects truth from two dangers at once: silence (nothing can be said) and idolatry (everything we say is identical to God).

Humility here is not epistemic surrender. It is epistemic hygiene. It keeps theology accountable by preventing the mind from confusing its constructs with the divine reality. In practice, analogical knowledge also makes moral space for correction, repentance, and growth. A theology that cannot be corrected tends to become a tool for power. A theology that is humble about its limits can still be firm about its commitments, but it remains open to refinement and purification.

So analogy does not dissolve truth; it disciplines it.

7.5 “This is special pleading”: why the thesis makes risky, falsifiable distinctions

A final objection says: the thesis is just moving goalposts. Whenever animals or machines show a human-like trait, you redefine the “real” human trait to stay ahead. That looks like special pleading: protecting a conclusion rather than following evidence.

The reply is to emphasize that the thesis is not based on a single “magic trait” that can be chased around. It is based on an integrated triad and on accountability-centered tests that are genuinely risky. For example, if a nonhuman entity demonstrated stable obligation-bearing agency under cost, persistent ownership of moral fault, capacity for repentance as turning rather than as output, and covenant-like bindingness that is not externally imposed, the thesis would face real pressure to revise its boundary claims. Those are not vague, untestable standards. They are demanding, but they are intelligible.

Likewise, if it turned out that strong normativity and repentance-like repair are robustly present in animal life in ways that go beyond social conditioning and into owned moral fault, the human uniqueness claim would have to be reformulated. The thesis is not immune to evidence; it is choosing the kinds of evidence that match the kind of claim being made.

Finally, the thesis takes a risk in the opposite direction: it refuses to treat output-fluency as decisive. In a culture trained to equate performance with reality, that refusal looks stubborn. But it is principled. It says: when categories are at stake, we must privilege accountability, obligation, and agency over surface similarity. That is not special pleading; it is a methodological choice aimed at protecting both human dignity and human responsibility in an era tempted to outsource both.

8. Christological Tightening: Image Restored and Clarified

A discussion of imago Dei can remain at the level of philosophical anthropology—capacity, normativity, and vocation—without explicitly naming Christ. But for Christianity, that would be an incomplete account. Christology "tightens" the doctrine by giving it a center of gravity. The image is not finally defined by our best abstract description of human faculties; it is defined by the One who is the image in fullness. That shift matters for two reasons. First, it protects the doctrine from becoming vague humanism ("humans are special because humans are special"). Second, it protects it from being reduced to cognitive achievement ("humans image God because humans are smart").

Christological tightening reframes imago Dei as both ontological and moral. Ontological: humans truly bear the image as creatures made for God-reference and moral address. Moral: the image can be distorted and can be restored. The image is not merely a static stamp; it is a living vocation ordered toward likeness. Christ becomes the measure not only of what the image is, but of what it is for.

8.1 Christ as the perfect image: derivative imaging vs full imaging

The core Christian move is to distinguish between derivative imaging and full imaging. Humans image God truly but finitely. Christ images God fully and perfectly. That difference immediately clarifies the "image vs model" problem. If Christ is the image in fullness, then the human task is not to build a complete conceptual model of God, but to receive and reflect God in creaturely form—through truth, holiness, and love—without pretending to contain the divine reality.

Derivative imaging means our likeness is real but partial. Humans can speak truth about God, but not exhaustively. Humans can manifest love, but not as pure love. Humans can practice justice, but always under judgment and in need of correction. That partiality is not a flaw to be embarrassed about; it is the natural condition of creaturehood.

Full imaging, in Christ, means there is a human life that is not merely God-referential but God-revealing. In Christ, the image is not an inference from human capacities; it is embodied disclosure. This is not a claim that "Jesus was the smartest." It is a claim that in Jesus, God is

made known in a way that is not reducible to theological language alone. The image becomes concrete: not a theory about what humans are, but a person who shows what God is like and what humanity is meant to be.

This helps the paper's argument in a subtle way. When machines generate impressive theology, it is tempting to treat theological correctness as the main measure of "image." Christological tightening resists that. The measure is not merely verbal correctness but lived fidelity. The "image" is not a clever description of God; it is likeness to God in being and action. That is why "derivative vs full" imaging matters: it prevents a drift into language-centered anthropology.

8.2 Conformity vs cleverness: truth as formation, not merely description

The modern world has trained people to equate truth with description: accurate statements, correct predictions, valid arguments. These are important. But Christian anthropology insists that truth is also formative. Truth is not only something we say; it is something we are conformed to.

This distinction matters because machines can become "clever" in the descriptive register. They can produce plausible theological explanations, generate biblically flavored counsel, and even simulate humility. But cleverness is not conformity. Conformity is the shaping of the self under moral and spiritual claims—repentance, obedience, patience, courage, love of enemy, mercy, faithfulness. Conformity includes cost. It is not a set of outputs; it is an orientation of the will over time.

In Christological terms, the image is clarified by imitation: humans do not merely theorize about God; they are called to become like Christ. That is why the image has a teleological direction. It is not a static stamp that guarantees righteousness. It is a gifted capacity and calling that demands growth. The truth of the doctrine is therefore tested not only in debates but in fruit—what kind of life it produces.

This also recalibrates how we interpret "God-talk." In Christian practice, prayer and worship are not primarily about informational content; they are practices of allegiance. They train the self. They turn the will. They are moral-spiritual acts that reshape the person toward love and humility. A machine that generates prayer language may help a human find words, but it does not itself undergo formation. It does not become holy. It does not learn obedience. It does not bear cost.

So Christological tightening provides an epistemic check: if we measure "image" by cleverness alone, we will accidentally grant image-status to any system that becomes verbally sophisticated. If we measure it by conformity—formation under truth—we preserve the distinction between simulation and moral agency.

8.3 Repair language: distortion, restoration, and moral transformation

Finally, Christological tightening introduces repair language. The image is not only given; it can be distorted. Distortion does not mean the image is erased; it means the image is misaligned. Human God-reference can become idolatry; moral normativity can become hypocrisy; vocation can become domination. The very capacities that make imago Dei meaningful can be turned into instruments of self-exaltation and harm.

This repair language helps the paper avoid idealizing humans. The argument is not "humans are uniquely good." It is that humans are uniquely accountable. The more gifted the capacity, the more devastating the distortion. Humans can weaponize concepts, justify cruelty with theology, and outsource responsibility to machines. The doctrine of distortion explains why: the image is real, but misdirected.

Restoration, then, is not primarily an upgrade in intelligence. It is moral transformation. It is the reordering of orientation, normativity, and vocation toward God. It includes repentance (turning), reconciliation (repair of relation), and sanctification (formation over time). In that sense, restoration is a repair of agency. It makes a person more truly a bearer of the image not by adding new faculties but by aligning existing faculties to their proper end.

This is crucial in the machine age. The temptation is to treat the crisis as cognitive: humans are being outperformed, so we need better skills. Christological repair says the deeper crisis is moral and spiritual: humans are tempted to abdicate responsibility, to treat tools as authorities, and to trade formation for fluency. Restoration is therefore a reassertion of human accountability and vocation: humans must remain the ones who repent, who discern, who bear cost, who seek truth with humility, and who steward power without domination.

Christological tightening, then, does not merely add religious ornament to the argument. It anchors the doctrine in a concrete center (Christ as the perfect image), redefines truth in formative terms (conformity, not cleverness), and explains both human failure and human hope through repair language (distortion and restoration). In a world dazzled by outputs, it insists that the image is most clearly seen not in what can be said, but in what can be lived—and transformed.

9. Implications for the Centaur Era: Humans With Powerful Tools

If the previous sections are right, then the defining challenge of the present moment is not that machines are becoming human, but that humans are tempted to become machine-like: to treat language as sufficient, outputs as authority, and responsibility as something that can be offloaded. The "centaur era" names a more mature framing: humans plus powerful tools,

intertwined. The question is not whether we will use these tools—we will—but whether we will use them in a way that preserves the image-bearing structure of human life: God-reference, addressability, accountability, and vocation.

The centaur posture is therefore neither technophobic nor credulous. It treats AI as real power, but as instrument power—amplification of human agency, not a transfer of moral agency from human to tool. In theological contexts, the stakes are unusually high, because the subject matter invites ultimacy claims, and ultimacy claims invite spiritual authority. A disciplined centaur practice aims to gain the benefits of amplification without surrendering the core human responsibilities that imago Dei entails.

9.1 AI as amplifier of description, not replacement of addressability

AI is exceptionally good at description: summarizing, paraphrasing, drafting, comparing arguments, retrieving patterns across large corpora, and generating candidate language quickly. In theology, that means it can help assemble outlines, propose interpretations, surface competing viewpoints, and draft prayers or liturgical text. This descriptive power is not trivial; it can raise the floor of literacy and make serious theological work more accessible.

But the imago Dei framework implies a limit: description is not addressability. A tool can help produce words about God, but it cannot be the moral-spiritual subject who stands under God's address. It cannot be the one who repents, submits, believes, or is transformed. Those are not merely cognitive actions; they are allegiance actions.

So the centaur claim is practical: use AI to amplify the descriptive layer—language, structure, synthesis—while explicitly reserving addressability for persons. Put bluntly: AI can help you talk, but it cannot take your place in the "You" relation to God or neighbor.

This distinction also protects human dignity. If we treat the most advanced descriptive capacity as the measure of personhood, then we risk elevating tools and devaluing humans whose descriptive capacities are limited. But if addressability remains primary, then the infant, the elderly, and the cognitively impaired remain fully within the circle of "you," regardless of performance. The centaur era must keep that circle intact.

9.2 Spiritual risks: outsourcing discernment and confusing fluency for authority

The danger in theological use is not mainly that AI will be "evil." The danger is that AI will be *convincing*. Fluency is spiritually hazardous because it feels like authority. A well-phrased answer can masquerade as wisdom; a confident tone can be mistaken for discernment; a coherent synthesis can be mistaken for truth.

This produces predictable failure modes.

First, outsourced discernment. People begin to treat the tool as a guide rather than as a drafting instrument. The tool's output becomes the end rather than the beginning of reflection. In that mode, prayer becomes text consumption, not address; moral deliberation becomes "what the model said," not accountable judgment.

Second, synthetic certainty. AI can generate a seamless narrative that resolves tensions too quickly. But theology often requires patience with mystery, careful attention to context, and willingness to remain under a humility ceiling. When a tool gives instant closure, it can train the soul toward premature certainty and away from reverent hesitation.

Third, authority laundering. Institutions may cite AI-generated theology to give an aura of neutrality or consensus: "the analysis shows…" This can become a new form of unaccountable authority, especially if users cannot audit inputs, data biases, or the chain of reasoning.

Fourth, idolatry-by-interface. Humans are relational creatures. If the AI speaks with empathy and spiritual vocabulary, people may bond to it as a quasi-pastoral presence. That bond can displace real community, real accountability, and real correction. It can also become spiritually deformative if the "counsel" is subtly self-serving, manipulative, or merely wrong.

These risks are not hypothetical. They are the predictable result of confusing linguistic resemblance with moral authority. The imago Dei framework warns: the more God-talk a tool can mimic, the more vigilant humans must be about where spiritual authority properly resides.

9.3 Practices: humility, accountability, and disciplined use in theological work

The solution is not abstinence; it is discipline. A centaur practice for theology should be built around the virtues implied by imago Dei: humility under God, accountability for one's words, and stewardship of power. Concretely, that means adopting practices that prevent the tool from quietly becoming an oracle.

A first practice is role clarity. Use AI as a drafting assistant, indexer, and summarizer—not as confessor, pastor, or spiritual director. Treat outputs as candidates for judgment, not as judgments. In personal spiritual work, keep the "address" portion—confession, prayer, repentance, discernment—strictly human.

A second practice is accountability anchoring. Every theological claim produced with AI assistance should have a human owner who is willing to say, "This is my claim; I stand behind it; I will revise it if wrong; I will bear the consequences if it harms." That single sentence disciplines the entire workflow.

A third practice is friction by design. Because AI reduces friction, it can also reduce discernment. Intentionally reintroduce friction at decision points: pause before publishing; ask what alternative interpretation a wise opponent would offer; consult primary sources; seek counsel from real humans; read aloud what you plan to teach. The goal is not slower work; it is truer work.

A fourth practice is confession of limits. In theological writing, it is legitimate to say, "This is analogical; I am not claiming comprehension; I may be wrong." That posture is not weakness; it is anti-idolatry. It also inoculates against the tool's tendency to sound certain.

A fifth practice is community tethering. Theology is not merely text; it is formation within a community. Use AI to help prepare, but keep interpretation accountable to real relationships: mentors, pastors, peers, and lived responsibilities. A tool cannot provide the kind of corrective feedback that comes from being known and being answerable.

These practices do not require technical sophistication. They require moral seriousness: a refusal to let convenience dissolve accountability.

9.4 Education: teaching "image vs model" as a literacy for the AI age

If the centaur era is permanent, we need a new basic literacy: the ability to distinguish an image from a model. This paper's earlier distinction becomes pedagogical: teach people to recognize when they are treating God as an object to be mastered (model) rather than as the One to whom they are accountable (image).

This literacy has several components.

First, genre awareness. Teach students that prayer, confession, lament, and worship are not merely textual genres. They are acts of allegiance. A machine can generate the genre, but it cannot perform the allegiance. This helps people use AI-generated drafts without mistaking them for spiritual life.

Second, authority awareness. Teach that fluency is not authority. Require students to trace claims to sources and to articulate what they personally affirm and why. Encourage the habit of slow reading and primary texts, especially in theological contexts where secondhand synthesis can be seductive.

Third, accountability training. Teach that moral and spiritual speech must have an owner. In an AI age, "I didn't write it" is not an excuse; it is a warning sign. Train people to sign their names only to what they can defend and to revise publicly when they err.

Fourth, formation over performance. Teach that the measure of theological work is not cleverness but fruit: does it produce humility, love, justice, truthfulness, courage, and care for

the vulnerable? This is a deeply Christian measure, and it is also the measure that resists AI-induced performative spirituality.

If this literacy becomes common, then the centaur era can be a period of amplification without idolatry: tools strengthen human expression while humans remain the morally accountable, spiritually addressable agents they were created to be. That is the practical implication of imago Dei for a world of persuasive machines.

10. Conclusion: Human Uniqueness Without Anti-Science or Machine Idolatry

The modern argument about mind has become unusually noisy because our reference points have shifted. Animals have proven more cognitively complex than simplistic hierarchies allowed, and machines have become more linguistically fluent than most people imagined possible. Under that pressure, many conversations collapse into a single axis—"intelligence"—and then swing between two extremes: denying animal depth to protect human dignity, or granting personhood to machines because they talk well. Both extremes are unstable. The first is empirically brittle and ethically callous. The second is spiritually dangerous and morally evasive.

This paper has argued for a third way: a disciplined account of imago Dei that respects science and resists idolatry. Human uniqueness, on this view, is not a brag about cognitive horsepower. It is a claim about what kind of agent a human is: a God-referential knower who can be addressed, held accountable, and tasked with stewardship. That framing can honor animal intelligence without collapse and can use machines powerfully without treating them as moral subjects. It also re-centers the truly urgent issue of the AI age: not whether tools can speak, but whether humans will remain answerable for what tools do in the world.

10.1 Restating the thesis: God-reference, addressability, accountability, vocation

The thesis can now be stated cleanly. Humans bear the image of God in a way that is structurally marked by four interlocking capacities.

God-reference is the capacity to conceptualize ultimacy without reducing God to an object—an orientation of thought and life toward the One who is not merely an explanatory hypothesis but a personal referent.

Addressability is the capacity to be hailed as a "you" morally and spiritually—summoned by claims that bind, corrected by truth, called by love, confronted by holiness.

Accountability is the capacity to own fault and to repair—guilt that is more than embarrassment, repentance that is more than output, promise-keeping and fidelity that are morally meaningful because an agent is truly answerable.

Vocation is the capacity to be entrusted—delegated dominion understood as stewardship, not domination; responsibility for neighbor, creation, and truth, including responsibility for the tools we build.

These four together mark a kind of agency that cannot be captured by a single "intelligence" scale. They describe a personhood grammar: what it means to stand before God and neighbor as a responsible subject rather than as a clever mechanism.

10.2 What is gained: clearer categories for animals, humans, and machines

A major gain of this framing is categorical clarity. It allows us to say several things at once without contradiction.

Animals can be cognitively impressive and morally significant without being human. Continuity is real; contempt is unwarranted. Animals deserve care, and human dominion must be accountable. But continuity does not force collapse into sameness, because the human signature includes strong normativity, explicit ultimacy-directed symbolism, and accountable agency in covenant-like forms.

Humans retain stable dignity without making dignity depend on performance. The image is not an IQ trophy. It is a creaturely status that grounds moral addressability across the whole range of human life, including the vulnerable. That stability matters precisely now, when societies are tempted to value persons by measurable outputs.

Machines can be extraordinarily useful without being moral subjects. We can recognize functional competence and treat it as genuine power while refusing the category error of assigning guilt, repentance, or spiritual authority to an instrument. This protects human accountability from evaporating into "the AI decided," and it protects spiritual life from being captivated by fluency mistaken for wisdom.

These clearer categories also reduce confusion in the "centaur era." Humans can use tools as amplifiers of description while preserving the non-delegable core of moral and spiritual agency: the "you" who can be called, corrected, forgiven, and tasked.

10.3 What remains open: hard edges, empirical humility, and ongoing dialogue

The argument has also left open what should remain open. The boundary claims here are not meant to be smug or closed to evidence. They are meant to be disciplined.

First, the hard edges are hard for a reason. It is genuinely difficult to specify, in purely empirical terms, where proto-normativity in animals ends and strong normativity begins, or what would count as unambiguous evidence of obligation-bearing agency in a nonhuman system. These are deep questions because they touch interiority and accountability, not merely behavior.

Second, empirical humility is required. Animal cognition research continues to evolve; simplistic dismissals are irresponsible. Likewise, AI systems continue to change; the landscape is not static. The paper's claim is about the kind of evidence needed for personhood and moral agency, not a prophecy that no surprises will occur.

Third, dialogue is necessary across disciplines. Theological anthropology must listen to ethology, cognitive science, and AI research, not to capitulate, but to avoid straw men. Conversely, scientific and technological communities should acknowledge that questions of obligation, accountability, and moral subjecthood are not reducible to performance benchmarks.

In short, the paper defends a boundary, but it does not deny that boundary-discernment is an ongoing task requiring both courage and humility.

10.4 Forward work: research questions and the need for a disciplined Reference section

Several research questions emerge naturally from this framework.

One set concerns animals: What are the strongest documented cases of proto-normativity, reconciliation, and rule-like behavior across species, and how should they be interpreted without either contempt or anthropomorphic inflation? How do symbolic capacities differ across species, and what is the relationship between symbolism and ultimacy-directed thought?

A second set concerns machines: What markers, beyond output, would credibly indicate obligation-bearing agency? How can we design AI systems and institutions so that responsibility remains traceable and cannot be laundered through "autonomy" rhetoric? What governance structures best prevent authority laundering in moral and spiritual domains?

A third set concerns formation: How do AI-mediated practices reshape human attention, conscience, prayer, and community? Does habitual reliance on fluent synthesis train the soul toward passivity and premature certainty, or can it be used to deepen discernment when disciplined by humility and accountability?

Finally, there is the scholarly task: a disciplined Reference section is not an afterthought; it is part of the paper's integrity. Because the argument spans theology, philosophy, cognitive science, and AI ethics, the references must be curated carefully: primary theological sources on imago Dei and analogy, major works on moral agency and personhood, credible surveys in animal cognition, and rigorous discussions of AI systems and accountability. The purpose of that section will not be to decorate the paper with citations, but to anchor claims to the best available work and to model the very posture this paper recommends: reverence plus audit, conviction under humility, and responsibility that does not hide behind fluency.

References

Youvan, Douglas C. *All Models Are Idols: Iconoclasm, Apophatic Humility, and the Refusal of Conceptual Possession* (preprint). 2026. doi:10.13140/RG.2.2.15429.10725.

Youvan, Douglas C. *Perception as Interface: Human and AI Misreadings of Reality* (preprint). 2025. doi:10.13140/RG.2.2.23412.41603.

Youvan, Douglas C., and GPT-5. *The Agnostic Manifold: GPT-5's Kolmogorovian τ and the Loss of Moral Directionality* (preprint). 2025. doi:10.13140/RG.2.2.35814.28483.

The Holy Bible. Genesis 1:26–28; Psalm 8; Colossians 1:15–20; Hebrews 1:1–4; Romans 8:29; 2 Corinthians 4:4. (Use the translation cited in the final manuscript.)

Aquinas, Thomas. *Summa Theologiae*. Translated by the Fathers of the English Dominican Province. Benziger Bros., 1947 (and later editions).

Aquinas, Thomas. *Summa Contra Gentiles*. Translated by Anton C. Pegis et al. University of Notre Dame Press, 1975.

Augustine. *On the Trinity (De Trinitate)*. Translated by Edmund Hill. New City Press, 1991.

Barth, Karl. *Church Dogmatics*. Vol. III/1 (The Doctrine of Creation). T&T Clark, 1958 (and later editions).

Calvin, John. *Institutes of the Christian Religion*. Translated by Ford Lewis Battles. Westminster John Knox Press, 1960.

Dionysius the Areopagite (Pseudo-Dionysius). *The Divine Names* and *Mystical Theology*. In *Pseudo-Dionysius: The Complete Works*, translated by Colm Luibheid. Paulist Press, 1987.

Irenaeus of Lyons. *Against Heresies*. Translated by Dominic J. Unger and John J. Dillon. Paulist Press, 1992.

Hoekema, Anthony A. *Created in God's Image*. Eerdmans, 1986.

Kilner, John F. *Dignity and Destiny: Humanity in the Image of God*. Eerdmans, 2015.

Middleton, J. Richard. *The Liberating Image: The Imago Dei in Genesis 1*. Brazos Press, 2005.

Kelsey, David H. *Eccentric Existence: A Theological Anthropology*. Westminster John Knox Press, 2009.

Pannenberg, Wolfhart. *Anthropology in Theological Perspective*. Westminster John Knox Press, 1985.

Przywara, Erich. *Analogia Entis: Metaphysics: Original Structure and Universal Rhythm*. Translated by John R. Betz and David Bentley Hart. Eerdmans, 2014.

Burrell, David B. *Knowing the Unknowable God: Ibn-Sina, Maimonides, Aquinas*. University of Notre Dame Press, 1986.

MacIntyre, Alasdair. *After Virtue: A Study in Moral Theory*. University of Notre Dame Press, 1981 (and later editions).

Taylor, Charles. *Sources of the Self: The Making of the Modern Identity*. Harvard University Press, 1989.

Smith, James K. A. *Desiring the Kingdom: Worship, Worldview, and Cultural Formation*. Baker Academic, 2009.

Buber, Martin. *I and Thou*. Translated by Walter Kaufmann. Scribner, 1970.

Levinas, Emmanuel. *Totality and Infinity: An Essay on Exteriority*. Translated by Alphonso Lingis. Duquesne University Press, 1969.

Kant, Immanuel. *Groundwork of the Metaphysics of Morals*. Translated by Mary Gregor. Cambridge University Press, 1998.

Korsgaard, Christine M. *The Sources of Normativity*. Cambridge University Press, 1996.

Korsgaard, Christine M. *Fellow Creatures: Our Obligations to the Other Animals*. Oxford University Press, 2018.

Tomasello, Michael. *A Natural History of Human Morality*. Harvard University Press, 2016.

de Waal, Frans. *Are We Smart Enough to Know How Smart Animals Are?* W. W. Norton, 2016.

de Waal, Frans. *Chimpanzee Politics: Power and Sex Among Apes*. Johns Hopkins University Press, 1982 (and later editions).

Shettleworth, Sara J. *Cognition, Evolution, and Behavior*. Oxford University Press, 2010.

Clayton, Nicola S., and Anthony Dickinson. “Episodic-like memory during cache recovery by scrub jays.” *Nature* 395 (1998): 272–274.

Singer, Peter. *Animal Liberation*. HarperCollins, 1975 (and later editions).

Turing, Alan M. “Computing Machinery and Intelligence.” *Mind* 59, no. 236 (1950): 433–460.

Searle, John R. “Minds, Brains, and Programs.” *Behavioral and Brain Sciences* 3, no. 3 (1980): 417–457.

Weizenbaum, Joseph. *Computer Power and Human Reason: From Judgment to Calculation*. W. H. Freeman, 1976.

Shannon, Claude E. “A Mathematical Theory of Communication.” *Bell System Technical Journal* 27 (1948): 379–423, 623–656.

Vaswani, Ashish, et al. “Attention Is All You Need.” In *Advances in Neural Information Processing Systems* 30 (2017).

Brown, Tom B., et al. “Language Models are Few-Shot Learners.” In *Advances in Neural Information Processing Systems* 33 (2020).

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” In *Proceedings of ACM FAccT* (2021).

OpenAI. *GPT-4 Technical Report*. 2023.

Russell, Stuart. *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking, 2019.

Floridi, Luciano. *The Fourth Revolution: How the Infosphere Is Reshaping Human Reality*. Oxford University Press, 2014.

Coeckelbergh, Mark. *AI Ethics*. MIT Press, 2020.

Crawford, Kate. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press, 2021.

O’Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown, 2016.

Diakopoulos, Nicholas. “Algorithmic Accountability: Journalistic Investigation of Computational Power Structures.” *Digital Journalism* 3, no. 3 (2015): 398–415.

Mitchell, Margaret, et al. “Model Cards for Model Reporting.” In *Proceedings of ACM FAT* (2019).

National Institute of Standards and Technology (NIST). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. 2023.

Kasparov, Garry. *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*. PublicAffairs, 2017.

Ellul, Jacques. *The Technological Society*. Translated by John Wilkinson. Knopf, 1964.

Postman, Neil. *Technopoly: The Surrender of Culture to Technology*. Knopf, 1992.

Borgmann, Albert. *Technology and the Character of Contemporary Life: A Philosophical Inquiry*. University of Chicago Press, 1984.

Heidegger, Martin. “The Question Concerning Technology.” In *The Question Concerning Technology and Other Essays*, translated by William Lovitt. Harper & Row, 1977.

The author has published over 4,700 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.